



Dual-Core Intel® Xeon® Processor 7000 Sequence

Specification Update

January 2006

Notice: The Dual-Core Intel® Xeon® processor 7000 sequence may contain design defects or errors known as errata that may cause the product to deviate from published specifications. Current characterized errata are documented in this specification update.

Document Number: 309627-003



INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY RELATING TO SALE AND/OR USE OF INTEL PRODUCTS, INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT, OR OTHER INTELLECTUAL PROPERTY RIGHT.

Intel products are not intended for use in medical, life saving, life sustaining, critical control or safety systems, or in nuclear facility applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an ordering number and are referenced in this document, or other Intel literature may be obtained by calling 1-800-548-4725 or by visiting Intel's website at <http://developer.intel.com/design/litcentr>.

Intel, Intel Xeon, Pentium, Pentium III, Celeron, and Intel NetBurst are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Copyright © 2005-2006, Intel Corporation. All rights reserved.

*Other names and brands may be claimed as the property of others.



Contents

Revision History	5
Preface.....	6
Identification Information	7
Summary Tables of Changes	9
Errata	14
Specification Changes.....	36
Specification Clarifications	37
Documentation Changes.....	38





Revision History

Version	Description	Date
-001	Initial release of the <i>Dual-Core Intel® Xeon® Processor 7000 Sequence Specification Update</i>	November 2005
-002	Added erratum A57.	December 2005
-003	Updated erratum A17; added errata A58-A79.	January 2006

Preface

This document is an update to the specifications contained in the Affected Documents and Related Documents tables below. This document is a compilation of device and documentation errata, specification clarifications and changes. It is intended for hardware system manufacturers and software developers of applications, operating systems, or tools.

Information types defined in Nomenclature are consolidated into the specification update and are no longer published in other documents.

This document may also contain information that was not previously published.

Affected Documents

1. *Dual-Core Intel® Xeon® Processor 7000 Sequence Datasheet*, Revision 1.0 (Document Number 309626)

Nomenclature

S-Spec Number is a five-digit code used to identify products. Products are differentiated by their unique characteristics, e.g., core speed, L3 cache size, package type, etc. as described in the processor identification information table. Read all notes associated with each S-Spec number.

Errata are design defects or errors. These may cause the processor's behavior to deviate from published specifications. Hardware and software designed to be used with any given stepping must assume that all errata documented for that stepping are present on all devices.

Specification Changes are modifications to the current published specifications. These changes will be incorporated in any new release of the specification.

Specification Clarifications describe a specification in greater detail or further highlight a specification's impact to a complex design situation. These clarifications will be incorporated in any new release of the specification.

Documentation Changes include typos, errors, or omissions from the current published specifications. These will be incorporated in any new release of the specification.

Note: Errata remain in the specification update throughout the product's lifecycle, or until a particular stepping is no longer commercially available. Under these circumstances, errata removed from the specification update are archived and available upon request. Specification changes, specification clarifications and documentation changes are removed from the specification update when the appropriate changes are made to the appropriate product specification or user documentation (datasheets, manuals, etc.).

Identification Information

Dual-Core Intel® Xeon® Processor 7000 Sequence Package Markings

Figure 1. Top-Side Marking Example





Figure 2. Bottom-Side Marking Example



The Dual-Core Intel Xeon processor 7000 sequence can be identified by the following register contents:

Family ¹	Model ²
1111b	0100b

NOTES:

1. The Family corresponds to bits [11:8] of the EDX register after RESET, bits [11:8] of the EAX register after the CUID instruction is executed with a 1 in the EAX register, and the generation field of the Device ID register accessible through Boundary Scan.
2. The Model corresponds to bits [7:4] of the EDX register after RESET, bits [7:4] of the EAX register after the CUID instruction is executed with a 1 in the EAX register, and the model field of the Device ID register accessible through Boundary Scan.

Cache and TLB descriptor parameters are provided in the EAX, EBX, ECX and EDX registers after the CUID instruction is executed with a 2 in the EAX register. Please refer to the *AP-485 Intel® Processor Identification and the CUID Instruction* Application Note and the *Prescott, Nocona and Potomac Processor BIOS Writer's Guide* for further information on the CUID instruction.

Table 1. Dual-Core Intel® Xeon® Processor 7000 Sequence Identification Information

QDF/ S-Spec	Core Stepping	L2 Cache Size (bytes)	CUID	Core Freq (GHz)	Data Bus Freq (MHz)	Package and Revision	Notes
SL8UC	A0	2M x 2	0F48h	3.0	667	604-pin micro-PGA with 53.3 x 53.3 mm FC-PGA4 package	1, 2, 3, 4, 5, 6, 7, 8
SL8UA	A0	1M x 2	0F48h	2.66	667	604-pin micro-PGA with 53.3 x 53.3 mm FC-PGA4 package	1, 2, 3, 4, 5, 6, 7, 8
SL8UD	A0	2M x 2	0F48h	3.0	800	604-pin micro-PGA with 53.3 x 53.3 mm FC-PGA4 package	1, 2, 3, 4, 5, 6, 7, 8
SL8UB	A0	1M x 2	0F48h	2.8	800	604-pin micro-PGA with 53.3 x 53.3 mm FC-PGA4 package	1, 2, 3, 4, 5, 6, 7, 8


NOTES:

1. These parts have PROCHOT# enabled
2. These parts have THERMTRIP# enabled
3. These parts are enabled for Enhanced Intel SpeedStep® Technology (EIST).
4. These parts are enabled for Enhanced Halt State (C1E).
5. These parts are enabled with Hyper-Threading Technology.
6. These parts are enabled with Execute Disable Bit (NX).
7. These parts are enabled for Intel® Extended Memory 64 Technology.
8. These parts are enabled with Intel® Virtualization Technology (VT).

Summary Tables of Changes

The following table indicates the Errata, Specification Changes, Specification Clarifications, or Documentation Changes which apply to the Dual-Core Intel® Xeon® processor 7000 sequence. Intel may fix some of the errata in a future stepping of the component, and account for the other outstanding issues through documentation or specification changes as noted. This table uses the following notations:

Codes Used in Summary Table

X:	Erratum, Specification Change or Clarification that applies to the given processor stepping.
(No mark) or (Blank Box):	This erratum is fixed in listed stepping or specification change does not apply to listed stepping.
Doc:	Document change or update that will be implemented.
Plan Fix:	This erratum may be fixed in a future of the product.
Fixed:	This erratum has been previously fixed.
No Fix:	There are no plans to fix this erratum.
	Change bar to left of table row indicates this item is either new or modified from the previous version of this document.

Each Specification Update item will be prefixed with a capital letter to distinguish the product. The key below details the letters that are used in Intel's microprocessor Specification Updates:

A	= Dual-Core Intel® Xeon® processor 7000 sequence
B	= Mobile Intel® Pentium® II processor
C	= Intel® Celeron® processor
D	= Dual-Core Intel® Xeon® processor 2.80 GHz
E	= Intel® Pentium® III processor
F	= Intel® Pentium® processor Extreme Edition and Intel® Pentium® D processor
G	= Intel® Pentium® III Xeon® processor
H	= Mobile Intel® Celeron® processor at 466/433/400/366/333/300 and 266 MHz
J	= 64-bit Intel® Xeon® processor MP with 1 MB L2 cache
K	= Mobile Intel® Pentium® III processor
L	= Intel® Celeron® D processor
M	= Mobile Intel® Celeron® processor
N	= Intel® Pentium® 4 processor
O	= Intel® Xeon® processor MP
P	= Intel® Xeon® processor
Q	= Mobile Intel® Pentium® 4 processor supporting Hyper-Threading Technology on 90-nm process technology
R	= Intel® Pentium® 4 processor on 90 nm process
S	= 64-bit Intel® Xeon® processor with 800 MHz system bus (1 MB and 2 MB L2 cache versions)
T	= Mobile Intel® Pentium® 4 processor-M

U = 64-bit Intel® Xeon® processor MP with up to 8 MB L3 cache
 V = Mobile Intel® Celeron® processor on .13 Micron Process in Micro-FCPGA Package
 W = Intel® Celeron® M processor
 X = Intel® Pentium® M processor on 90nm process with 2-MB L2 Cache
 Y = Intel® Pentium® M processor
 Z = Mobile Intel® Pentium® 4 processor with 533 MHz system bus
 AC = Intel® Celeron® processor in 678-pin Package

The Specification Updates for the Pentium® processor, Pentium® Pro processor, and other Intel products do not use this convention.

Errata (Sheet 1 of 3)

No.	A0	Plans	Description
A1	X	No Fix	Transaction is not retired after BINIT#
A2	X	No Fix	Invalid opcode 0FFFh requires a modrm byte
A3	X	No Fix	Processor may hang due to speculative page walks to non-existent system memory
A4	X	No Fix	Memory type of the load lock different from its corresponding store unlock
A5	X	No Fix	Machine check architecture error reporting and recovery may not work as expected
A6	X	No Fix	Debug mechanisms may not function as expected
A7	X	No Fix	Cascading of performance counters does not work correctly when forced overflow is enabled
A8	X	No Fix	EMON event counting of x87 loads may not work as expected
A9	X	No Fix	System bus interrupt messages without data and which receive a hard-failure response may hang the processor
A10	X	No Fix	The processor signals page fault exception (#PF) instead of alignment check exception (#AC) on an unlocked CMPXCHG8B instruction
A11	X	No Fix	FSW may not be completely restored after page fault on FRSTOR or FLDDENV instructions
A12	X	No Fix	Processor issues inconsistent transaction size attributes for locked operation
A13	X	No Fix	When the processor is in the system management mode (SMM), debug registers may be fully writeable
A14	X	No Fix	Shutdown and IERR# may result due to a machine check exception on a Hyper-Threading Technology enabled processor
A15	X	No Fix	Processor may hang under certain frequencies and 12.5% STPCLK# duty cycle
A16	X	No Fix	System may hang if a fatal cache error causes bus write line (BWL) transaction to occur to the same cache line address as an outstanding bus read line (BRL) or bus read-invalidate line (BRIL)
A17	X	No Fix	A write to an APIC register sometimes may appear to have not occurred
A18	X	No Fix	Parity error in the L1 cache may cause the processor to hang
A19	X	No Fix	Locks and SMC detection may cause the processor to temporarily hang
A20	X	No Fix	Incorrect debug exception (#DB) may occur when a data breakpoint is set on an FP instruction
A21	X	No Fix	xAPIC may not report some illegal vector error
A22	X	No Fix	Incorrect duty cycle is chosen when on-demand clock modulation is enabled in a processor supporting Hyper-Threading Technology
A23	X	No Fix	Memory aliasing of pages as uncacheable memory type and write back (WB) may hang the system
A24	X	No Fix	Interactions between the instruction translation lookaside buffer (ITLB) and the instruction streaming buffer may cause unpredictable software behavior
A25	X	No Fix	Using STPCLK# and executing code from very slow memory could lead to a system hang
A26	X	No Fix	Processor provides a 4-byte store unlock after an 8-Byte load lock

Errata (Sheet 2 of 3)

No.	A0	Plans	Description
A27	X	No Fix	Data breakpoints on the high half of a floating point line split may not be captured
A28	X	No Fix	Machine Check exceptions may not update last-exception record MSRs (LERs)
A29	X	No Fix	MOV CR3 performs incorrect reserved bit checking when in PAE paging
A30	X	No Fix	Stores to page tables may not be visible to pagewalks for subsequent loads without serializing or invalidating the page table entry
A31	X	No Fix	Processor may fault when the upper 8 bytes of segment selector is loaded from a far jump through a call gate via the local descriptor table
A32	X	No Fix	Loading a stack segment with a selector that references a non-canonical address can lead to a #SS fault on a processor supporting Intel® Extended Memory 64 Technology (Intel® EM64T)
A33	X	No Fix	FXRSTOR may not restore non-canonical effective addresses on processors with Intel® Extended Memory 64 Technology (Intel® EM64T) Enabled
A34	X	No Fix	A push of esp that faults may zero the upper 32 bits of RSP
A35	X	No Fix	Enhanced halt state (C1E) may not be entered in a Hyper-Threading Technology enabled processor
A36	X	No Fix	Checking of page table base address may not match the address bit width supported by the platform
A37	X	No Fix	IA32_MCi_STATUS MSR may improperly indicate that additional MCA information may have been captured
A38	X	No Fix	With trap flag (TF) asserted, FP instruction that triggers an unmasked FP exception may take single step trap before retirement of instruction
A39	X	No Fix	Branch trace store (BTS) and precise event based sampling (PEBS) may update memory outside the BTS/PEBS buffer
A40	X	No Fix	Memory ordering failure may occur with snoop filtering third party agents after issuing and completing a bus write invalidate line (BWIL) or bus locked write (BLW) transaction
A41	X	No Fix	Control register 2 (CR2) can be updated during a REP MOVS/STOS instruction with fast strings enabled
A42	X	No Fix	REP STOS/MOVS instructions with RCX >= 2 ³² may cause a system hang
A43	X	No Fix	An REP MOVS or an REP STOS instruction with RCX >= 2 ³² may fail to execute to completion or may write to incorrect memory locations on processors supporting Intel® Extended Memory 64 Technology (Intel® EM64T)
A44	X	No Fix	An REP LODSB or an REP LODSD or an REP LODSQ instruction with RCX >= 2 ³² may cause a system hang on processors supporting Intel® Extended Memory 64 Technology (Intel® EM64T)
A45	X	No Fix	Data access which spans both canonical and non-canonical address space may hang system
A46	X	No Fix	Running in SMM and L1 data cache adaptive mode may cause unexpected system behavior when SMRAM is mapped to cacheable memory
A47	X	No Fix	A 64-bit value of linear instruction pointer (LIP) may be reported incorrectly in the branch trace store (BTS) memory record or in the precise event based sampling (PEBS) memory record
A48	X	No Fix	PDE/PTE Loads and continuous locked updates to the same cache line may cause a system livelock
A49	X	No Fix	At core-to-bus ratios of 16:1 and above defer reply transactions with non-zero REQb Values; may cause a front side bus stall
A50	X	No Fix	CPUID reports thermal monitor 2 supported when running at ratios 18:1 and above
A51	X	No Fix	The processor may issue front side bus transactions up to 6 clocks after RESET# is asserted
A52	X	No Fix	Front side bus machine checks may be reported as a result of on-going transactions during warm reset
A53	X	No Fix	Entering single logical processor mode via power on configuration may not work
A54	X	No Fix	Machine check exception may be signaled in a system with multiple threads and several lock transactions
A55	X	No Fix	The processor may issue multiple code fetches to the same cache line for systems with slow memory

Errata (Sheet 3 of 3)

No.	A0	Plans	Description
A56	X	No Fix	Writing the local vector table (LVT) when an interrupt is pending may cause an unexpected interrupt
A57	X	No Fix	IRET under certain conditions may cause an unexpected alignment check exception
A58	X	Plan Fix	Running with Virtual Machine Extensions (VMX) in L1 data cache adaptive mode may cause unexpected system behavior
A59	X	Plan Fix	A mispredicted branch may issue a speculative load to an incorrect address during VM exit on processors supporting Intel® Virtualization Technology
A60	X	Plan Fix	Exit qualification and pending debug exceptions Virtual-Machine Control Structure (VMCS) fields contain incorrect information on VM exits due to debug exceptions
A61	X	Plan Fix	VM exit saves incorrect interruptibility state information on exit due to nested exceptions
A62	X	Plan Fix	VM exit on Load Machine Status Word (LMSW) may not show expected exit information in the Virtual-Machine Control Structure (VMCS)
A63	X	Plan Fix	An incorrect load may be issued under conditions which cause VM exit
A64	X	Plan Fix	A Start-up IPI (SIPI) VM exit does not clear pending INITs
A65	X	No Fix	Access to an unsupported address range in uniprocessor (UP) or dual processor (DP) systems supporting Intel® Virtualization Technology may not trigger appropriate actions
A66	X	No Fix	VM exit due to a MOV from CR8 May cause an unexpected memory access
A67	X	No Fix	The processor may incorrectly respond to machine checks during VM entry/exit transitions
A68	X	No Fix	NIT during string operations in the Virtual-Machine Extension (VMX) guest mode may cause unexpected system behavior
A69	X	No Fix	Power down requests may not be serviced if a power down transition is interrupted by an in-target probe event in the presence of a specific type of VM exit
A70	X	No Fix	VM entry/exit writes to LSTAR/SYSCALL_FLAG MSRs may cause incorrect data to be written to bits [63:32]
A71	X	No Fix	VM exit due to TPR shadow below threshold may improperly set and cause "Blocking by STI" actions
A72	X	Plan Fix	VM exit on Load Machine Status Word (LMSW) may not show expected exit information in the Virtual-Machine Control Structure (VMCS)
A73	X	Plan Fix	A VM exit may occur when the processor is in wait-for-SIPI or shutdown states and a chipwide power down transition occurs
A74	X	No Fix	The execution of a VMPTRLD instruction may cause an unexpected memory access
A75	X	No Fix	The execution of VMPTRLD or VMREAD may cause an unexpected memory access
A76	X	Plan Fix	Attempting to use an LDT entry when the LDTR has been loaded with an unusable segment may cause unexpected memory accesses
A77	X	No Fix	FS/GS base MSRs can be loaded from MSR-load areas during VM entry or VM exit
A78	X	Plan Fix	NMI-blocking information recorded in VMCS may be incorrect after a #GP on an IRET instruction
A79	X	Plan Fix	VMLAUNCH/VMRESUME may not fail when VMCS is programmed to cause VM exit to return to a different mode

Specification Changes

No.	SPECIFICATION CHANGES
	None for this revision of this specification update.

Specification Clarifications

No.	SPECIFICATION CLARIFICATIONS
	None for this revision of this specification update.

Documentation Changes

No.	DOCUMENTATION CHANGES
	None for this revision of this specification update.

Errata

A1. Transaction is not retired after BINIT#

Problem: If the first transaction of a locked sequence receives a HITM# and DEFER# during the snoop phase it should be retried and the locked sequence restarted. However, if BINIT# is also asserted during this transaction, the transaction will not be retried.

Implication: When this erratum occurs, locked transactions will not be retried.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A2. Invalid opcode 0FFFh requires a modrm byte

Problem: Some invalid opcodes require a ModRM byte and other following bytes, while others do not. The invalid opcode 0FFFh did not require a ModRM in previous generation microprocessors such as Pentium® II or Pentium III processors, but it is required in the Intel® Xeon® processor.

Implication: The use of an invalid opcode 0FFFh without the ModRM byte may result in a page or limit fault on the Intel Xeon processor. When this erratum occurs, locked transactions will not be retried.

Workaround: To avoid this erratum use ModRM byte with invalid 0FFFh opcode.

Status: For the steppings affected, see the *Summary Table of Changes*.

A3. Processor may hang due to speculative page walks to non-existent system memory

Problem: A load operation that misses the data translation lookaside buffer (DTLB) will result in a page walk. If the page walk loads the page directory entry (PDE) from cacheable memory and that PDE load returns data that points to a valid page table entry (PTE) in uncacheable memory the processor will access the address referenced by the PTE. If the address referenced does not exist the processor will hang with no response from system memory.

Implication: Processor may hang due to speculative page walks to non-existent system memory.

Workaround: Page directories and page tables in UC memory space must point to system memory that exists.

Status: For the steppings affected, see the *Summary Table of Changes*.

A4. Memory type of the load lock different from its corresponding store unlock

Problem: The Intel Xeon Processor employs a use-once protocol to ensure that a processor in a multiprocessor system may access data that are loaded into its cache on a Read-for-Ownership operation at least once before it is snooped out by another processor. This protocol is necessary to avoid a dual processor livelock scenario where no processor in the system can gain ownership of a line and modify it before those data are snooped out by another processor. In the case of this erratum, the use-once protocol incorrectly activates for split load lock instructions. A load lock operation accesses data that split across a page boundary with both pages of WB memory type. The use-once protocol activates and the memory type for the split halves get forced to UC. Since use-once does not apply to stores, the store unlock instructions go out as WB memory type. The full sequence on the Bus is: locked partial read (UC), partial read (UC), partial write (WB), locked partial write (WB). The Use-once protocol should not be applied to Load locks.

Implication: When this erratum occurs, the memory type of the load lock will be different than the memory type of the store unlock operation. This behavior (Load Locks and Store Unlocks having different

memory types) does not however introduce any functional failures such as system hangs or memory corruption.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A5. Machine check architecture error reporting and recovery may not work as expected

Problem: When the processor detects errors it should attempt to report and/or recover from the error. In the situations described below, the processor does not report and/or recover from the error(s) as intended.

- When a transaction is deferred during the snoop phase and subsequently receives a Hard Failure response, the transaction should be removed from the bus queue so that the processor may proceed. Instead, the transaction is not properly removed from the bus queue, the bus queue is blocked, and the processor will hang.
- When a hardware prefetch results in an uncorrectable tag error in the L2 cache, MC0_STATUS.UNCOR and MC0_STATUS.PCC are set but no Machine Check Exception (MCE) is signaled. No data loss or corruption occurs because the data being prefetched has not been used. If the data location with the uncorrectable tag error is subsequently accessed, an MCE will occur. However, upon this MCE, or any other subsequent MCE, the information for that error will not be logged because MC0_STATUS.UNCOR has already been set and the MCA status registers will not contain information about the error which caused the MCE assertion but instead will contain information about the prefetch error event.
- When the reporting of errors is disabled for Machine Check Architecture (MCA) Bank 2 by setting all MC2_CTL register bits to 0, uncorrectable errors should be logged in the IA32_MC2_STATUS register but no machine-check exception should be generated. Uncorrectable loads on bank 2, which would normally be logged in the IA32_MC2_STATUS register, are not logged.
- When one half of a 64 byte instruction fetch from the L2 cache has an uncorrectable error and the other 32 byte half of the same fetch from the L2 cache has a correctable error, the processor will attempt to correct the correctable error but cannot proceed due to the uncorrectable error. When this occurs the processor will hang.
- When an L1 cache parity error occurs, the cache controller logic should write the physical address of the data memory location that produced that error into the IA32_MC1_ADDR REGISTER (MC1_ADDR). In some instances of a parity error on a load operation that hits the L1 cache, however, the cache controller logic may write the physical address from a subsequent load or store operation into the IA32_MC1_ADDR register.
- When an error exists in the tag field of a cache line such that a request for ownership (RFO) issued by the processor hits multiple tag fields in the L2 cache (the correct tag and the tag with the error) and the accessed data information also has a correctable error, the processor will correctly log the multiple tag match error but will hang when attempting to execute the machine check exception handler.
- If a memory access receives a machine check error on both 64 byte halves of a 128-byte L2 cache sector, the IA32_MC0_STATUS register records this event as multiple errors, i.e., the valid error bit and the overflow error bit are both set indicating that a machine check error occurred while the results of a previous error were in the error-reporting bank. The IA32_MC1_STATUS register should also record this event as multiple errors but instead records this event as only one correctable error.
- The overflow bit should be set to indicate when more than one error has occurred. The overflow bit being set indicates that more than one error has occurred. Because of this erratum,

if any further errors occur, the MCA overflow bit will not be updated, thereby incorrectly indicating only one error has been received.

- If an I/O instruction (IN, INS, REP INS, OUT, OUTS, or REP OUTS) is being executed, and if the data for this instruction become corrupted, the processor will signal a Machine Check Exception (MCE). If the instruction is directed at a device that is powered down, the processor may also receive an assertion of SMI#. Since MCEs have higher priority, the processor will call the MCE handler, and the SMI# assertion will remain pending. However, while attempting to execute the first instruction of the MCE handler, the SMI# will be recognized and the processor will attempt to execute the SMM handler. If the SMM handler is successfully completed, it will attempt to restart the I/O instruction, but will not have the correct machine state due to the call to the MCE handler. This can lead to failure of the restart and shutdown of the processor.
- If PWRGOOD is de-asserted during a RESET# assertion causing internal glitches, the MCA registers may latch invalid information.
- If RESET# is asserted, then de-asserted, and reasserted, before the processor has cleared the MCA registers, then the information in the MCA registers may not be reliable, regardless of the state or state transitions of PWRGOOD.
- If MCERR# is asserted by one processor and observed by another processor, the observing processor does not log the assertion of MCERR#. The Machine Check Exception (MCE) handler called upon assertion of MCERR# will not have any way to determine the cause of the MCE.
- The Overflow Error bit (bit 62) in the IA32_MC0_STATUS register indicates, when set, that a machine check error occurred while the results of a previous error were still in the error reporting bank (i.e. The Valid bit was set when the new error occurred). If an uncorrectable error is logged in the error-reporting bank and another error occurs, the overflow bit will not be set.
- The MCA Error Code field of the IA32_MC0_STATUS register gets written by a different mechanism than the rest of the register. For uncorrectable errors, the other fields in the IA32_MC0_STATUS register are only updated by the first error. Any further errors that are detected will update the MCA Error Code field without updating the rest of the register, thereby leaving the IA32_MC0_STATUS register with stale information.
- When a speculative load operation hits the L2 cache and receives a correctable error, the IA32_MC1_Status Register may be updated with incorrect information. The IA32_MC1_Status Register should not be updated for speculative loads.
- The processor should only log the address for L1 parity errors in the IA32_MC1_Status register if a valid address is available. If a valid address is not available, the Address Valid bit in the IA32_MC1_Status register should not be set. In instances where an L1 parity error occurs and the address is not available because the linear to physical address translation is not complete or an internal resource conflict has occurred, the Address Valid bit is incorrectly set.
- The processor may hang when an instruction code fetch receives a hard failure response from the Front Side Bus. This occurs because the bus control logic does not return data to the core, leaving the processor empty. IA32_MC0_STATUS MSR does indicate that a hard fail response occurred.

The processor may hang when the following events occur and the machine check exception is enabled, CR4.MCE=1. A processor that has its STPCLK# pin asserted will internally enter the Stop Grant State and finally issue a Stop Grant Acknowledge special cycle to the bus. If an uncorrectable error is generated during the Stop Grant process it is possible for the Stop Grant special cycle to be issued to the bus before the processor vectors to the machine check handler.

Once the chipset receives its last Stop Grant special cycle it is allowed to ignore any bus activity from the processors. As a result, processor accesses to the machine check handler may not be acknowledged, resulting in a processor hang.

Implication: The processor is unable to correctly report and/or recover from certain errors

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A6. Debug mechanisms may not function as expected

Problem: If the first transaction of a locked sequence receives a HITM# and DEFER# during the snoop phase it should be retried and the locked sequence restarted. However, if BINIT# is also asserted during this transaction, the transaction will not be Certain debug mechanisms may not function as expected on the processor. The cases are as follows:

- When the following conditions occur: 1) An FLD instruction signals a stack overflow or underflow, 2) the FLD instruction splits a page-boundary or a 64 byte cache line boundary, 3) the instruction matches a Debug Register on the high page or cache line respectively, and 4) the FLD has a stack fault and a memory fault on a split access, the processor will only signal the stack fault and the debug exception will not be taken.
- When a data breakpoint is set on the ninth and/or tenth byte(s) of a floating point store using the Extended Real data type, and an unmasked floating point exception occurs on the store, the break point will not be captured.
- When any instruction has multiple debug register matches, and any one of those debug registers is enabled in DR7, all of the matches should be reported in DR6 when the processor goes to the debug handler. This is not true during a REP instruction. As an example, during execution of a REP MOVSW instruction the first iteration a load matches DR0 and DR2 and sets DR6 as FFFF0FF5h. On a subsequent iteration of the instruction, a load matches only DR0. The DR6 register is expected to still contain FFFF0FF5h, but the processor will update DR6 to FFFF0FF1h.

A data breakpoint that is set on a load to uncacheable memory may be ignored due to an internal segment register access conflict. In this case the system will continue to execute instructions, bypassing the intended breakpoint. Avoiding having instructions that access segment descriptor registers e.g. LGDT, LIDT close to the UC load, and avoiding serialized instructions before the UC load will reduce the occurrence of this erratum.

Implication: Certain debug mechanisms do not function as expected on the processor.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A7. Cascading of performance counters does not work correctly when forced overflow is enabled

Problem: The performance counters are organized into pairs. When the CASCADE bit of the Counter Configuration Control Register (CCCR) is set, a counter that overflows will continue to count in the other counter of the pair. The FORCE_OVF bit forces the counters to overflow on every non-zero increment. When the FORCE_OVF bit is set, the counter overflow bit will be set but the counter no longer cascades.

Implication: The performance counters do not cascade when the FORCE_OVF bit is set.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A8. EMON event counting of x87 loads may not work as expected

Problem: If a performance counter is set to count x87 loads and floating-point exceptions are unmasked, the FPU operand (Data) pointer (FDP) may become corrupted.

Implication: When this erratum occurs, FPU operand (Data) pointer (FDP) may become corrupted.

Workaround: This erratum will not occur with floating point exceptions masked. If floating-point exceptions are unmasked, then performance counting of x87 loads should be disabled.

Status: For the steppings affected, see the *Summary Table of Changes*.

A9. System bus interrupt messages without data and which receive a hard-failure response may hang the processor

Problem: When a System Bus agent (processor or chipset) issues an interrupt transaction without data onto the System Bus, and the transaction receives a hard-failure response, a potential processor hang can occur. The processor, which generates an inter-processor interrupt (IPI) that receives hard-failure response, will still log the MCA error event cause as hard-failure, even if the APIC causes a hang. Other processors, which are true targets of the IPI, will also hang on hard-failure-without-data, but will not record an MCA hard-failure event as a cause. If a hard-failure response occurs on a System Bus interrupt message with data, the APIC will complete the operation so as not to hang the processor.

Implication: The processor may hang.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A10. The processor signals page fault exception (#PF) instead of alignment check exception (#AC) on an unlocked CMPXCHG8B instruction

Problem: If a page fault exception (#PF) and alignment check exception (#AC) both occur for an unlocked CMPXCHG8B instruction, then #PF will be flagged.

Implication: Software that depends on the (#AC) before the (#PF) will be affected since #PF is signaled in this case.

Workaround: Remove the software's dependency on #AC having precedence over #PF. Alternately, correct the page fault in the page fault handler and then restart the faulting instruction.

Status: For the steppings affected, see the *Summary Table of Changes*.

A11. FSW may not be completely restored after page fault on FRSTOR or FLDENV instructions

Problem: If the FPU operating environment or FPU state (operating environment and register stack) being loaded by an FLDENV or FRSTOR instruction wraps around a 64-Kbyte or 4-Gbyte boundary and a #PF or segment limit fault (#GP or #SS) occurs on the instruction near the wrap boundary, the upper byte of the FPU status word (FSW) might not be restored. If the fault handler does not restart program execution at the faulting instruction, stale data may exist in the FSW.

Implication: When this erratum occurs, stale data will exist in the FSW.

Workaround: Ensure that the FPU operating environment and FPU state do not cross 64-Kbyte or 4-Gbyte boundaries. Alternately, ensure that the page fault handler restarts program execution at the faulting instruction after correcting the paging problem.

Status: For the steppings affected, see the *Summary Table of Changes*.

A12. Processor issues inconsistent transaction size attributes for locked operation

Problem: When the processor is in the page address extension (PAE) mode and detects the need to set the Access and/or Dirty bits in the page directory or page table entries, the processor sends an 8 byte load lock onto the System Bus. A subsequent 8 byte store unlock is expected, but instead a 4 byte store unlock occurs. Correct data are provided since only the lower bytes change, however external logic monitoring the data transfer may be expecting an 8-byte store unlock.

Implication: No known commercially available chipsets are affected by this erratum.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A13. When the processor is in the system management mode (SMM), debug registers may be fully writeable

Problem: When in system management mode (SMM), the processor executes code and stores data in the SMRAM space. When the processor is in this mode and writes are made to DR6 and DR7, the processor should block writes to the reserved bit locations. Due to this erratum, the processor may not block these writes. This may result in invalid data in the reserved bit locations.

Implication: Reserved bit locations within DR6 and DR7 may become invalid.

Workaround: Software may perform a read/modify/write when writing to DR6 and DR7 to ensure that the values in the reserved bits are maintained.

Status: For the steppings affected, see the *Summary Table of Changes*.

A14. Shutdown and IERR# may result due to a machine check exception on a Hyper-Threading Technology enabled processor

Problem: When a machine check exception (MCE) occurs due to an internal error, both logical processors on a Hyper-Threading Technology enabled processor normally vector to the MCE handler. However, if one of the logical processors is in the “Wait for SIPI” state, that logical processor will not have a MCE handler and will shut down and assert IERR#.

Implication: A processor with a logical processor in the “Wait for SIPI” state will shut down when an MCE occurs on the other thread.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A15. Processor may hang under certain frequencies and 12.5% STPCLK# duty cycle

Problem: If a system de-asserts STPCLK# at a 12.5% duty cycle, and the processor is running below 2 GHz, and the processor thermal control circuit (TCC) on-demand clock modulation is active, the processor may hang. This erratum does not occur under the automatic mode of the TCC.

Implication: When this erratum occurs, the processor will hang.

Workaround: If use of the on-demand mode of the processor's TCC is desired in conjunction with STPCLK# modulation, then assure that STPCLK# is not asserted at a 12.5% duty cycle.

Status: For the steppings affected, see the *Summary Table of Changes*.

A16. System may hang if a fatal cache error causes bus write line (BWL) transaction to occur to the same cache line address as an outstanding bus read line (BRL) or bus read-invalidate line (BRIL)

Problem: A processor internal cache fatal data ECC error may cause the processor to issue a BWL transaction to the same cache line address as an outstanding BRL or BRIL. As it is not typical

behavior for a single processor to have a BWL and a BRL/BRIL concurrently outstanding to the same address, this may represent an unexpected scenario to system logic within the chipset.

Implication: The processor may not be able to fully execute the machine check handler in response to the fatal cache error if system logic does not ensure forward progress on the System Bus under this scenario.

Workaround: System logic should ensure completion of the outstanding transactions. Note that during recovery from a fatal data ECC error, memory image coherency of the BWL with respect to BRL/BRIL transactions is not important. Forward progress is the primary requirement.

Status: For the steppings affected, see the *Summary Table of Changes*.

A17. A write to an APIC register sometimes may appear to have not occurred

Problem: With respect to the retirement of instructions, stores to the uncacheable memory-based APIC register space are handled in a non-synchronized way. For example if an instruction that masks the interrupt flag, e.g. CLI, is executed soon after an uncacheable write to the Task Priority Register (TPR) that lowers the APIC priority, the interrupt masking operation may take effect before the actual priority has been lowered. This may cause interrupts whose priority is lower than the initial TPR, but higher than the final TPR, to not be serviced until the interrupt enabled flag is finally set, i.e. by STI instruction. Interrupts will remain pending and are not lost.

Implication: In this example the processor may allow interrupts to be accepted but may delay their service.

Workaround: This non-synchronization can be avoided by issuing an APIC register read after the APIC register write. This will force the store to the APIC register before any subsequent instructions are executed. No commercial operating system is known to be impacted by this erratum.

Status: For the steppings affected, see the *Summary Table of Changes*.

A18. Parity error in the L1 cache may cause the processor to hang

Problem: If a locked operation accesses a line in the L1 cache that has a parity error, it is possible that the processor may hang while trying to evict the line.

Implication: If this erratum occurs, it may result in a system hang. Intel has not observed this erratum with any commercially available software.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A19. Locks and SMC detection may cause the processor to temporarily hang

Problem: The processor may temporarily hang in an HT Technology enabled system, if one logical processor executes a synchronization loop that includes one or more bus locks and is waiting for release by the other logical processor. If the releasing logical processor is executing instructions that are within the detection range of the self modifying code (SMC) logic, then the processor may be locked in the synchronization loop until the arrival of an interrupt or other event.

Implication: If this erratum occurs in an HT Technology enabled system, the application may temporarily stop making forward progress. Intel has not observed this erratum with any commercially available software.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A20. Incorrect debug exception (#DB) may occur when a data breakpoint is set on an FP instruction

Problem: The default Microcode Floating Point Event Handler routine executes a series of loads to obtain data about the FP instruction that are causing the FP event. If a data breakpoint is set on the

instruction causing the FP event, the load in the microcode routine will trigger the data breakpoint resulting in a Debug Exception.

Implication: An incorrect Debug Exception (#DB) may occur if data breakpoint is placed on an FP instruction. Intel has not observed this erratum with any commercially available software or system.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A21. xAPIC may not report some illegal vector error

Problem: The local xAPIC has an Error Status Register, which records all errors it detects. Bit 6 of this register, the Receive Illegal Vector bit, is set when the local xAPIC detects an illegal vector in a message that it receives. When an illegal vector error is received on the same internal clock that the error status register is being written due to a previous error, bit 6 does not get set and illegal vector errors are not flagged.

Implication: The xAPIC may not report some Illegal Vector errors when they occur at approximately the same time as other xAPIC errors. The other xAPIC errors will continue to be reported.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A22. Incorrect duty cycle is chosen when on-demand clock modulation is enabled in a processor supporting Hyper-Threading Technology

Problem: When a processor supporting Hyper-Threading Technology enables On-Demand Clock Modulation on both logical processors, the processor is expected to select the lowest duty cycle of the two potentially different values. When one logical processor enters the AUTOHALT state, the duty cycle implemented should be unaffected by the halted logical processor. Due to this erratum, the duty cycle is incorrectly chosen to be the higher duty cycle of both logical processors.

Implication: Due to this erratum, higher duty cycle may be chosen when the On-Demand Clock Modulation is enabled on both logical processors.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A23. Memory aliasing of pages as uncachable memory type and write back (WB) may hang the system

Problem: When a page is being accessed as either Uncacheable (UC) or Write Combining (WC) and WB, under certain bus and memory timing conditions, the system may loop in a continual sequence of UC fetch, implicit writeback, and Request For Ownership (RFO) retries

Implication: This erratum has not been observed in any commercially available operating system or application. The aliasing of memory regions, a condition necessary for this erratum to occur, is documented as being unsupported in the *IA-32 Intel® Architecture Software Developer's Manual*, Volume 3, section 10.12.4, Programming the PAT. However, if this erratum occurs the system may hang

Workaround: The pages should not be mapped as either UC or WC and WB at the same time.

Status: For the steppings affected, see the *Summary Table of Changes*.

A24. Interactions between the instruction translation lookaside buffer (ITLB) and the instruction streaming buffer may cause unpredictable software behavior

Problem: Complex interactions within the instruction fetch/decode unit may make it possible for the processor to execute instructions from an internal streaming buffer containing stale or incorrect information.

Implication: When this erratum occurs, an incorrect instruction stream may be executed resulting in unpredictable software behavior.

Workaround: It is possible for the BIOS to contain a workaround for this erratum.

Status: For the steppings affected, see the *Summary Table of Changes*.

A25. Using STPCLK# and executing code from very slow memory could lead to a system hang

Problem: The system may hang when the following conditions are met:

1. Periodic STPCLK# mechanism is enabled via the chipset
2. Hyper-Threading Technology is enabled
3. One logical processor is waiting for an event (i.e. hardware interrupt)
4. The other logical processor executes code from very slow memory such that every code fetch is deferred long enough for the STPCLK# to be re-asserted.

Implication: If this erratum occurs, the processor will go into and out of the sleep state without making forward progress, since the logical processor will not be able to service any pending event. This erratum has not been observed in any commercial platform running commercial software.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A26. Processor provides a 4-byte store unlock after an 8-Byte load lock

Problem: When the processor is in the Page Address Extension (PAE) mode and detects the need to set the Access and/or Dirty bits in the page directory or page table entries, the processor sends an 8 byte load lock onto the system bus. A subsequent 8 byte store unlock is expected, but instead a 4 byte store unlock occurs. Correct data information is provided since only the lower bytes change, however external logic monitoring the data transfer may be expecting an 8 byte load lock.

Implication: No known commercially available chipsets are affected by this erratum.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A27. Data breakpoints on the high half of a floating point line split may not be captured

Problem: When a floating point load which splits a 64-byte cache line gets a floating point stack fault, and a data breakpoint register maps to the high line of the floating point load, internal boundary conditions exist that may prevent the data breakpoint from being captured.

Implication: When this erratum occurs, a data breakpoint will not be captured.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A28. Machine Check exceptions may not update last-exception record MSRs (LERs)

Problem: The Last-Exception Record MSRs (LERs) may not get updated when Machine Check Exceptions occur.

Implication: When this erratum occurs, the LER may not contain information relating to the machine check exception. They will contain information relating to the exception prior to the machine check exception.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A29. MOV CR3 performs incorrect reserved bit checking when in PAE paging

Problem: The MOV CR3 instruction should perform reserved bit checking on the upper unimplemented address bits. This checking range should match the address width reported by CPUID instruction 0x8000008. This erratum applies whenever PAE is enabled.

Implication: Software that sets the upper address bits on a MOV CR3 instruction and expects a fault may fail. This erratum has not been observed with commercially available software.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A30. Stores to page tables may not be visible to pagewalks for subsequent loads without serializing or invalidating the page table entry

Problem: Under rare timing circumstances, a page table load on behalf of a programmatically younger memory access may not get data from a programmatically older store to the page table entry if there is not a fencing operation or page translation invalidate operation between the store and the younger memory access. Refer to the *IA-32 Intel® Architecture Software Developer's Manual* for the correct way to update page tables. Software that conforms to the Software Developer's Manual will operate correctly.

Implication: If the guidelines in the Software Developer's Manual are not followed, stale data may be loaded into the processor's translation lookaside buffer (TLB) and used for memory operations. This erratum has not been observed with any commercially available software.

Workaround: The guidelines in the *IA-32 Intel® Architecture Software Developer's Manual* should be followed.

Status: For the steppings affected, see the *Summary Table of Changes*.

A31. Processor may fault when the upper 8 bytes of segment selector is loaded from a far jump through a call gate via the local descriptor table

Problem: In IA-32e mode of the Intel® EM64T processor, control transfers through a call gate via the local descriptor table (LDT) that uses a 16-byte descriptor, the upper 8-byte access may wrap and access an incorrect descriptor in the LDT. This only occurs on an LDT with a LIMIT > 0x10008 with a 16-byte descriptor that has a selector of 0xFFFFC.

Implication: In the event this erratum occurs, the upper 8-byte access may wrap and access an incorrect descriptor within the LDT, potentially resulting in a fault or system hang. Intel has not observed this erratum with any commercially available software.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A32. Loading a stack segment with a selector that references a non-canonical address can lead to a #SS fault on a processor supporting Intel® Extended Memory 64 Technology (Intel® EM64T)

Problem: When a processor supporting Intel EM64T is in IA-32e mode, loading a stack segment with a selector which references a non-canonical address will result in a #SS fault instead of a #GP fault.

Implication: When this erratum occurs, Intel EM64T enabled systems may encounter unexpected behavior.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A33. FXRSTOR may not restore non-canonical effective addresses on processors with Intel® Extended Memory 64 Technology (Intel® EM64T) Enabled

Problem: If an x87 data instruction has been executed with a non-canonical effective address, FXSAVE may store that non-canonical FP data pointer (FDP) value into the save image. An FXRSTOR instruction executed with 64-bit operand size may signal a General Protection Fault (#GP) if the FDP or FP instruction pointer (FIP) is in non-canonical form.

Implication: When this erratum occurs, Intel EM64T enabled systems may encounter an unintended #GP fault.

Workaround: Software should avoid using non-canonical effective addressing in Intel EM64T enabled processors. BIOS can contain a workaround for this erratum removing the unintended #GP fault on FXRSTOR.

Status: For the steppings affected, see the *Summary Table of Changes*.

A34. A push of esp that faults may zero the upper 32 bits of RSP

Problem: In the event that a push ESP instruction, that faults, is executed in compatibility mode, the processor will incorrectly zero upper 32-bits of RSP.

Implication: A Push of ESP in compatibility mode will zero the upper 32-bits of RSP. Due to this erratum, this instruction fault may change the contents of RSP. This erratum has not been observed in commercially available software.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A35. Enhanced halt state (C1E) may not be entered in a Hyper-Threading Technology enabled processor

Problem: If the IA32_MISC_ENABLE MSR (0x1A0) C1E enable bit is not set prior to an INIT event on an HT Technology enabled system, the processor will not enter C1E until the next SIPI wakeup event for the second logical processor.

Implication: Due to this erratum, the processor will not enter C1E state.

Workaround: If C1E is supported in the system, the IA32_MISC_ENABLE MSR should be enabled prior to issuing the first SIPI to the second logical processor.

Status: For the steppings affected, see the *Summary Table of Changes*.

A36. Checking of page table base address may not match the address bit width supported by the platform

Problem: If the page table base address, included in the page map level-4 table, page-directory pointer table, page-directory table or page table, exceeds the physical address range supported by the platform (e.g. 36-bit) and it is less than the implemented address range (e.g. 40-bit), the processor does not check if the address is invalid.

Implication: If software sets such invalid physical address in those tables, the processor does not generate a page fault (#PF) upon access to that virtual address, and the access results in an incorrect read or write. If BIOS provides only valid physical address ranges to the operating system, this erratum will not occur.

Workaround: BIOS must provide valid physical address ranges to the operating system.

Status: For the steppings affected, see the *Summary Table of Changes*.

A37. IA32_MCi_STATUS MSR may improperly indicate that additional MCA information may have been captured

Problem: When a data parity error is detected and the bus queue is busy, the ADDR_V and MISCV bits of the IA32_MCi_STATUS register may be asserted even though the contents of the IA32_MCi_ADDR and IA32_MCi_MISC MSRs were not properly captured.

Implication: If this erratum occurs, the MCA information captured in the IA32_MCi_ADDR and IA32_MCi_MISC may not correspond to the reported machine-check error, even though the ADDR_V and MISCV are asserted.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A38. With trap flag (TF) asserted, FP instruction that triggers an unmasked FP exception may take single step trap before retirement of instruction

Problem: If an FP instruction generates an unmasked exception with the EFLAGS.TF=1, it is possible for external events to occur, including a transition to a lower power state. When resuming from the lower power state, it may be possible to take the single step trap before the execution of the original FP instruction completes.

Implication: A Single Step trap will be taken when not expected.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A39. Branch trace store (BTS) and precise event based sampling (PEBS) may update memory outside the BTS/PEBS buffer

Problem: If the BTS/PEBS buffer is defined such that:

- The difference between BTS/PEBS buffer base and BTS/PEBS absolute maximum is not an integer multiple of the corresponding record sizes
- BTS/PEBS absolute maximum is less than a record size from the end of the virtual address space
- The record that would cross BTS/PEBS absolute maximum will also continue past the end of the virtual address space

A BTS/PEBS record can be written that will wrap at the 4G boundary (IA-32) or 2⁶⁴ boundary (Intel EM64T mode), and write memory outside of the BTS/PEBS buffer.

Implication: Software that uses BTS/PEBS near the 4 G boundary (IA-32) or 2⁶⁴ boundary (Intel EM64T mode), and defines the buffer such that it does not hold an integer multiple of records can update memory outside the BTS/PEBS buffer.

Workaround: Define BTS/PEBS buffer such that BTS/PEBS absolute maximum minus BTS/PEBS buffer base is integer multiple of the corresponding record sizes as recommended in the *IA-32 Intel® Architecture Software Developer's Manual*, Volume 3.

Status: For the steppings affected, see the *Summary Table of Changes*.

A40. Memory ordering failure may occur with snoop filtering third party agents after issuing and completing a bus write invalidate line (BWIL) or bus locked write (BLW) transaction

Problem: Under limited circumstances, the processors may, after issuing and completing a BWIL or BLW transaction, retain data from the addressed cache line in shared state even though the specification requires complete invalidation. This data retention may also occur when a BWIL transaction's self-snooping yields HITM snoop results.

Implication: A system may suffer memory ordering failures if its central agent incorporates coherence sequencing which depends on full self-invalidation of the cache line associated (1) with BWIL and BLW transactions, or (2) all HITM snoop results without regard to the transaction type and snoop results source.

Workaround: 1. The central agent can issue a bus cycle that causes a cache line to be invalidated (Bus Read Invalidate Line (BRIL) or BWIL transaction) in response to a processor-generated BWIL (or BLW) transaction to insure complete invalidation of the associated cache line. If there are no intervening processor-originated transactions to that cache line, the central agent's invalidating snoop will get a clean snoop result.

Or

2. Snoop filtering central agents can:
 - a. Not use processor-originated BWIL or BLW transactions to update their snoop filter information, or
 - b. Update the associated cache line state information to shared state on the originating bus (rather than invalid state) in reaction to a BWIL or BLW.

Status: For the steppings affected, see the *Summary Table of Changes*.

A41. Control register 2 (CR2) can be updated during a REP MOVS/STOS instruction with fast strings enabled

Problem: Under limited circumstances while executing a REP MOVS/STOS string instruction, with fast strings enabled, it is possible for the value in CR2 to be changed as a result of an interim paging event, normally invisible to the user. Any higher priority architectural event that arrives and is handled while the interim paging event is occurring may see the modified value of CR2.

Implication: The value in CR2 is correct at the time that an architectural page fault is signaled. Intel has not observed this erratum with any commercially available software.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A42. REP STOS/MOVS instructions with RCX >= 2³² may cause a system hang

Problem: In IA-32e mode using Intel EM64T-enabled processors, executing a repeating string instruction with the iteration count greater than or equal to 2³² and a pending event may cause the REP STOS/MOVS instruction to live lock and hang.

Implication: When this erratum occurs, the processor may live lock and result in a system hang. Intel has not observed this erratum with any commercially available software or system.

Workaround: Do not use strings larger than 4 GB.

Status: For the steppings affected, see the *Summary Table of Changes*.

A43. An REP MOVS or an REP STOS instruction with RCX >= 2³² may fail to execute to completion or may write to incorrect memory locations on processors supporting Intel® Extended Memory 64 Technology (Intel® EM64T)

Problem: In IA-32e mode using Intel EM64T-enabled processors, an REP MOVS or an REP STOS instruction executed with the register RCX >= 2³², may fail to execute to completion or may write data to incorrect memory locations.

Implication: This erratum may cause an incomplete instruction execution or incorrect data in the memory. Intel has not observed this erratum with any commercially available software or system.

Workaround: It is possible for the BIOS to contain a workaround for this erratum.

Status: For the steppings affected, see the *Summary Table of Changes*.

A44. An REP LODSB or an REP LODSD or an REP LODSQ instruction with RCX >= 2³² may cause a system hang on processors supporting Intel® Extended Memory 64 Technology (Intel® EM64T)

Problem: In IA-32e mode using Intel EM64T-enabled processors, an REP LODSB or an REP LODSD or an REP LODSQ instruction executed with the register RCX >= 2³² may fail to complete execution causing a system hang. Additionally, there may be no #GP fault due to the non-canonical address in the RSI register.

Implication: This erratum may cause a system hang on Intel EM64T-enabled platforms. Intel has not observed this erratum with any commercially available software or system.

Workaround: It is possible for the BIOS to contain a workaround for this erratum

Status: For the steppings affected, see the *Summary Table of Changes*.

A45. Data access which spans both canonical and non-canonical address space may hang system

Problem: If a data access causes a page split across the canonical to non-canonical address space, the processor may livelock which in turn would cause a system hang.

Implication: When this erratum occurs, the processor may livelock, resulting in a system hang. Intel has not observed this erratum with any commercially available software.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A46. Running in SMM and L1 data cache adaptive mode may cause unexpected system behavior when SMRAM is mapped to cacheable memory

Problem: In a Hyper-Threading Technology-enabled system, unexpected system behavior may occur if a change is made to the value of the CR3 result from an RSM (Resume from System Management) instruction while in L1 data cache adaptive mode (IA32_MISC_ENABLES MSR 0x1a0, bit 24). This behavior will only be visible when SMRAM is mapped into WB/WT cacheable memory on SMM entry and exit.

Implication: This erratum can have multiple failure symptoms, including incorrect data in memory. Intel has not observed this erratum with any commercially available software.

Workaround: Disable L1 data cache adaptive mode by setting the L1 Data Cache Context Mode control bit (bit 24) of the IA32_MISC_ENABLES MSR (0x1a0) to 1. It is possible for the BIOS to contain a workaround for this erratum

Status: For the steppings affected, see the *Summary Table of Changes*.

A47. A 64-bit value of linear instruction pointer (LIP) may be reported incorrectly in the branch trace store (BTS) memory record or in the precise event based sampling (PEBS) memory record

Problem: On a processors supporting Intel EM64T:

- If an instruction fetch wraps around the 4G boundary in Compatibility Mode, the 64-bit value of LIP in the BTS memory record will be incorrect (upper 32 bits will be set to 0xFFFFFFFF when they should be 0).
- If a PEBS event occurs on an instruction whose last byte is at memory location FFFFFFFFh, the 64-bit value of LIP in the PEBS record will be incorrect (upper 32 bits will be set to FFFFFFFFh when they should be 0).

Implication: Intel has not observed this erratum on any commercially available software.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A48. PDE/PTE Loads and continuous locked updates to the same cache line may cause a system livelock

Problem: In a multiprocessor configuration, if one processor is continuously doing locked updates to a cache line that is being accessed by another processor doing a page table walk, the page table walk may not complete.

Implication: Due to this erratum, the system may livelock until some external event interrupts the locked update. Intel has not observed this erratum with any commercially available software.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A49. At core-to-bus ratios of 16:1 and above defer reply transactions with non-zero REQb Values; may cause a front side bus stall

Problem: Certain processors are likely to hang the front side bus (FSB) if the following conditions are met:

1. A Defer Reply transaction has a REQb[2:0] value of either 010b, 011b, 100b, 110b, or 111b, and
2. The operating bus ratio is 16:1 or higher.

When these conditions are met, the processor may incorrectly and indefinitely assert a snoop stall for the Defer Reply transaction. Such an event will block further progress on the FSB.

Implication: If this erratum occurs, the system may hang. Intel has not observed this erratum with any commercially available system.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A50. CPUID reports thermal monitor 2 supported when running at ratios 18:1 and above

Problem: When the CPUID instruction is executed on a processor which is running at ratios 18:1 and above, the system incorrectly reports that Thermal Monitor 2 is supported.

Implication: When this erratum occurs, the processor incorrectly reports that Thermal Monitor 2 is supported.

Workaround: Software should ignore the feature flag on the processor and not use it as an indication that it can enable Thermal Monitor 2. Note that Thermal Monitor remains a feature and must be enabled for the processor to remain within specification.

Status: For the steppings affected, see the *Summary Table of Changes*.

A51. The processor may issue front side bus transactions up to 6 clocks after RESET# is asserted

Problem: The processor may issue transactions beyond the documented 3 FSB clocks and up to 6 FSB clocks after RESET# is asserted in the case of a warm reset. A warm reset is where the chipset asserts RESET# when the system is running.

Implication: The processor may issue transactions up to 6 FSB clocks after RESET# is asserted.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A52. Front side bus machine checks may be reported as a result of on-going transactions during warm reset

- Problem:** Processor FSB protocol/signal integrity machine checks may be reported if the transactions are initiated or in-progress during a warm reset. A warm reset is where the chipset asserts RESET# when the system is running.
- Implication:** The processor may log FSB protocol/signal integrity machine checks if transactions are allowed to occur during RESET# assertions.
- Workaround:** BIOS may clear FSB protocol/signal integrity machine checks for systems/chipsets which do not block new transactions during RESET# assertions.
- Status:** For the steppings affected, see the *Summary Table of Changes*.

A53. Entering single logical processor mode via power on configuration may not work

- Problem:** When the system uses power on configuration (POC) to enter single logical processor mode on a dual core processor (by asserting A31# at the deassertion of RESET#), the system may be susceptible to a variety of failing symptoms including; system hangs and MCERR# or IERR# assertions.
- Implication:** POC can not be used to enter single logical processor mode.
- Workaround:** None identified.
- Status:** For the steppings affected, see the *Summary Table of Changes*.

A54. Machine check exception may be signaled in a system with multiple threads and several lock transactions

- Problem:** The processor assumes the central agent is responsible for forward progress once it has retried a request. The central agent attempts to guarantee forward progress to each bus agent by periodically providing a slot where the requester can drive a transaction that will not get retried. The central agent does not track progress on a per thread basis. When the following conditions are met, it is possible for one thread to not make forward progress and result in a fatal machine check exception.
1. One agent performs repeated bus lock requests.
 2. Another agent performs repeated cache locks and also performs a code fetch for the other thread.
- Implication:** Due to this erratum, the processor may generate a fatal Machine Check Exception with a time-out exception error code.
- Workaround:** It is possible for the BIOS to contain a workaround for this erratum.
- Status:** For the steppings affected, see the *Summary Table of Changes*.

A55. The processor may issue multiple code fetches to the same cache line for systems with slow memory

- Problem:** Systems with long latencies on returning code fetch data from memory e.g., BIOS ROM, may cause the processor to issue multiple fetches to the same cache line, once per each instruction executed.
- Implication:** This erratum may slow down system boot time. Intel has not observed a failure, as a result of this erratum, in a commercially available system.
- Workaround:** None identified.
- Status:** For the steppings affected, see the *Summary Table of Changes*.

A56. Writing the local vector table (LVT) when an interrupt is pending may cause an unexpected interrupt

Problem: If a local interrupt is pending when the LVT entry is written, an interrupt may be taken on the new interrupt vector even if the mask bit is set.

Implication: An interrupt may immediately be generated with the new vector when a LVT entry is written, even if the new LVT entry has the mask bit set. If there is no interrupt service routine (ISR) set up for that vector the system will GP fault. If the ISR does not do an end of interrupt (EOI) the bit for the vector will be left set in the in-service register and mask all interrupts at the same or lower priority.

Workaround: Any vector programmed into an LVT entry must have an ISR associated with it, even if that vector was programmed as masked. This ISR routine must do an EOI to clear any unexpected interrupts that may occur. The ISR associated with the spurious vector does not generate an EOI, therefore the spurious vector should not be used when writing the LVT.

Status: For the steppings affected, see the *Summary Table of Changes*.

A57. IRET under certain conditions may cause an unexpected alignment check exception

Problem: In IA-32e mode, it is possible to get an alignment check exception (#AC) on the IRET instruction even though alignment checks were disabled at the start of the IRET. This can only occur if the IRET instruction is returning from CPL3 code to CPL3 code. IRETs from CPL0/1/2 are not affected. This erratum can occur if the EFLAGS value on the stack has the AC flag set, and the interrupt handler's stack is misaligned. In IA-32e mode, RSP is aligned to a 16-byte boundary before pushing the stack frame.

Implication: In IA-32e mode, under the conditions given above, an IRET can get a #AC even if alignment checks are disabled at the start of the IRET. This erratum can only be observed with a software generated stack frame.

Workaround: Software should not generate misaligned stack frames for use with IRET.

Status: For the steppings affected, see the *Summary Table of Changes*.

A58. Running with Virtual Machine Extensions (VMX) in L1 data cache adaptive mode may cause unexpected system behavior

Problem: In a system supporting Hyper-Threading Technology and Intel® Virtualization Technology, unexpected system behavior may result if a change to the value of CR0, CR3 and CR4 occurs during a VM entry or VM exit operation while in L1 data cache adaptive mode (IA32_MISC_ENABLER MSR 0x1a0 bit 24).

Implication: This erratum can have multiple failure symptoms including incorrect data in memory.

Workaround: It is possible for the BIOS to contain a workaround for this erratum.

Status: For the steppings affected, see the *Summary Table of Changes*.

A59. A mispredicted branch may issue a speculative load to an incorrect address during VM exit on processors supporting Intel® Virtualization Technology

Problem: During VM exit on processors supporting Intel Virtualization Technology, a mispredicted branch may result in a speculative load to an incorrect address.

Implication: This erratum may cause access to an incorrect address space during VM exit.

Workaround: It is possible for BIOS to have a workaround for this erratum.

Status: For the steppings affected, see the *Summary Table of Changes*.

A60. Exit qualification and pending debug exceptions Virtual-Machine Control Structure (VMCS) fields contain incorrect information on VM exits due to debug exceptions

Problem: In systems supporting Intel Virtualization Technology, a debug exception which causes a VM exit may cause incorrect information to be loaded into the pending debug exceptions and the Exit Qualification fields of the Virtual-Machine Control Structure (VMCS).

Implication: When executing code using debug breakpoints, host software may see incorrect information in these fields of the VMCS.

Workaround: It is possible for the BIOS to contain a workaround for this erratum.

Status: For the steppings affected, see the *Summary Table of Changes*.

A61. VM exit saves incorrect interruptibility state information on exit due to nested exceptions

Problem: VM exit does not save Interruptibility State field Bits[1:0] as 0 on an exit due to nested exception or double fault or triple fault. The VM exit may save spurious information about events that are being blocked by STI and MOVSS/POPSS instructions.

Implication: The Interruptibility State field Bits[1:0] may not contain appropriate state of guest events being blocked by STI or MOVSS/POPSS instruction.

Workaround: It is possible for the BIOS to contain a workaround for this erratum.

Status: For the steppings affected, see the *Summary Table of Changes*.

A62. VM exit on Load Machine Status Word (LMSW) may not show expected exit information in the Virtual-Machine Control Structure (VMCS)

Problem: In systems supporting Intel Virtualization Technology, a VM exit on an LMSW instruction, which references memory, may not return expected exit information concerning data breakpoints in the VMCS.

Implication: When this erratum occurs, Virtual-Machine Monitor (VMM) software may not be able to recognize and handle data breakpoints for the LMSW instruction

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A63. An incorrect load may be issued under conditions which cause VM exit

Problem: In systems supporting Intel Virtualization Technology, any condition which causes VM exit may cause the processor to issue a load to an invalid address if the Virtual Machine Extension (VMX) operation is configured to run with legacy treatment of SMI Delivery.

Implication: This erratum may cause unpredictable system behavior or system hang.

Workaround: It is possible for the BIOS to contain a workaround for this erratum.

Status: For the steppings affected, see the *Summary Table of Changes*.

A64. A Start-up IPI (SIPI) VM exit does not clear pending INITs

Problem: In systems supporting Intel Virtualization Technology, a SIPI initiated VM exit does not clear pending INITs.

Implication: Pending INITs will be serviced upon the return to the guest.

Workaround: It is possible for the BIOS to contain a workaround for this erratum.

Status: For the steppings affected, see the *Summary Table of Changes*.

A65. Access to an unsupported address range in uniprocessor (UP) or dual processor (DP) systems supporting Intel® Virtualization Technology may not trigger appropriate actions

Problem: When using processors supporting Intel Virtualization Technology and configured as dual or single processor-capable (i.e. not multiprocessor-capable), the processor should perform address checks using a maximum physical address width of 36. Instead, these processors will perform address checks using a maximum physical address width of 40.

Implication: Due to this erratum, actions which are normally taken upon detection of an unsupported address may not occur. Software which does not attempt to access unsupported addresses will not experience this erratum.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A66. VM exit due to a MOV from CR8 may cause an unexpected memory access

Problem: In a system supporting Intel Virtualization Technology and Intel® Extended Memory 64 Technology (Intel® EM64T), if the “CR8-store exiting” bit in the processor-based VM-execution control field is set and the “use TPR shadow” bit is not set, a MOV from CR8 instruction executed by a Virtual Machine Extensions (VMX) guest that causes a VM exit may generate an unexpected memory access.

Implication: When this erratum occurs, a read access to unexpected address may be issued to the chipset. Subsequent side effects are dependent on chipset operation and may include system hang.

Workaround: It is possible for the BIOS to contain a workaround for this erratum.

Status: For the steppings affected, see the *Summary Table of Changes*.

A67. The processor may incorrectly respond to machine checks during VM entry/exit transitions

Problem: In systems supporting Intel Virtualization Technology, when machine checks are encountered during VM entry/exit transitions, the processor is expected to respond with a VM exit (if a machine check occurs during VM entry) or abort (if a machine check occurs during VM exit). As a result of this erratum when machine checks occur during VM entry/exit transitions the processor will attempt to service the machine check which may lead to IERR-shutdown or execution of the Machine Check handler, dependent on the CR4.MCE setting.

Implication: The system may end up in the shutdown state if CR4.MCE is not set.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A68. NIT during string operations in the Virtual-Machine Extension (VMX) guest mode may cause unexpected system behavior

Problem: In a system supporting Intel Virtualization Technology, if INIT occurs during REP LODS/MOVS/STOS/INS/OUTS while the processor is executing in VMX guest mode, after servicing the INIT, the host will resume at the next instruction and does not complete the remainder of string operation.

Implication: This erratum may cause unexpected system behavior to occur.

Workaround: It is possible for the BIOS to contain a workaround for this erratum.

Status: For the steppings affected, see the *Summary Table of Changes*.

A69. Power down requests may not be serviced if a power down transition is interrupted by an in-target probe event in the presence of a specific type of VM exit

Problem: In a system supporting Intel Virtualization Technology, the processor may service a pended VM exit prior to completely exiting out of a low power state when the following sequences of events occur:

- Chip-wide power down transition occurs, and
- VM exit due to a VMLaunch, VMResume, STI, POPF, POPFD, or IRET instruction is pended, and
- Chip-wide power down transition is interrupted by an in-target probe event.

Implication: Due to this erratum the processor may not recognize further STPCLK# assertions, TM1, TM2, or Enhanced Intel SpeedStep® Technology. Intel has not observed this erratum with any commercially available software.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A70. VM entry/exit writes to LSTAR/SYSCALL_FLAG MSRs may cause incorrect data to be written to bits [63:32]

Problem: Incorrect MSR data in bits [63:32] may be observed in the following two cases:

1. When ECX contains 0xC0000084 and a VM entry/exit writes the IA32_CR_LSTAR MSR (MSR Address 0xC0000082) bits [63:32] of the data may be zeroed.
2. When ECX does not contain 0xC0000084 and a VM entry/exit writes the IA32_CR_SYSCALL_FLAG_MASK MSR (MSR Address 0xC0000084) bits [63:32] of the data may not be zeroed.

Implication: Bits [63:32] of the affected MSRs may contain the wrong data after a VM exit/entry which loads the affected MSR.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A71. VM exit due to TPR shadow below threshold may improperly set and cause “Blocking by STI” actions

Problem: In a system supporting Intel Virtualization Technology and Intel EM64T, the “blocking by STI” bit of the interruptibility-state field may be saved as 1 rather than 0. This erratum may occur when a STI instruction is executed directly prior to a MOV to CR8 which results in a VM exit due to a reduction of the TPR shadow value below the TPR threshold.

Implication: When this erratum occurs, delivery of an interrupt may be delayed by one instruction.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A72. VM exit on Load Machine Status Word (LMSW) may not show expected exit information in the Virtual-Machine Control Structure (VMCS)

Problem: In systems supporting Intel Virtualization Technology, a VM exit on an LMSW instruction, which references memory, may not return expected exit information concerning data breakpoints in the VMCS.

Implication: When this erratum occurs, VMM software may not be able to recognize and handle data breakpoints for the LMSW instruction.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A73. A VM exit may occur when the processor is in wait-for-SIPI or shutdown states and a chipwide power down transition occurs

Problem: In a system supporting Intel Virtualization Technology, the processor may incorrectly VM exit under the following conditions:

- Interrupt-Window-Exiting VM-execution control is set
- RFLAGS[IF]=1
- Chipwide power down transition requests occur when the processor is in Wait-For-SIPI or Shutdown states.

Implication: Due to this erratum, Interrupt-Window-Exiting VM exits may take the logical processor out of Wait-For-SIPI and Shutdown states. Intel has not observed this erratum with any commercially available software.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A74. The execution of a VMPTRLD instruction may cause an unexpected memory access

Problem: In a system supporting Intel Virtualization Technology, executing VMPTRLD may cause a memory access to an address not referenced by the memory operand.

Implication: This erratum may cause unpredictable system behavior or system hang.

Workaround: It is possible for the BIOS to contain a workaround for this erratum.

Status: For the steppings affected, see the *Summary Table of Changes*.

A75. The execution of VMPTRLD or VMREAD may cause an unexpected memory access

Problem: On processors supporting Intel Virtualization Technology, executing a VMPTRLD or a VMREAD instruction outside of VMX mode may result in a load to an unexpected address.

Implication: This erratum may cause a load to an unexpected memory address.

Workaround: It is possible for the BIOS to contain a workaround for this erratum.

Status: For the steppings affected, see the *Summary Table of Changes*.

A76. Attempting to use an LDT entry when the LDTR has been loaded with an unusable segment may cause unexpected memory accesses

Problem: In a system supporting Intel® EM64T and Intel Virtualization Technology when the following occur:

- The LDTR is loaded during VM entry with the segment unusable bit set for the LDTR in the VMCS.
- The segment limit is non-zero.
- The granularity bit is set references to a segment located in the LDT in 64-bit mode at any time later may cause the processor to exhibit unexpected behavior.

Implication: This erratum may cause unexpected memory accesses.

Workaround: None identified.

Status: For the steppings affected, see the *Summary Table of Changes*.

A77. FS/GS base MSR can be loaded from MSR-load areas during VM entry or VM exit

Problem: If the VM exit or VM entry MSR-load area contains references to the FS or GS Base MSRs, the VM exit and VM entry transitions should fail. Instead, the operation will load the MSRs with the value in the corresponding MSR-load area entry.

Implication: VM entries and VM exits that should fail will complete successfully in this situation. If a VM entry is to virtual-8086 mode, the base address for FS or for GS may be loaded with a value that is not consistent with that mode. Intel has not observed this erratum with any commercially available software or systems.

Workaround: Software should not enter values in the MSR-load areas that correspond to either the FS base MSR or the GS base MSR. Software can establish the value of these registers on VM entry using the guest-state area of the VMCS and on VM exit using the host-state area of the VMCS.

Status: For the steppings affected, see the *Summary Table of Changes*.

A78. NMI-blocking information recorded in VMCS may be incorrect after a #GP on an IRET instruction

Problem: In a system supporting Intel Virtualization Technology, the NMI blocking bit in the Interruption-Information field in the guest VMCS may be set incorrectly. This erratum will happen if a VM exit occurs for a #GP fault on an IRET instruction due to an EIP that violates the segment limit or is non-canonical.

Implication: If this erratum occurs, monitor software may not be able to handle #GP and then inject an NMI since monitor software does not have information about whether NMIs are blocked in the guest.

Workaround: Monitor software can workaround this bug by avoiding injection of NMI after #GP emulation.

Status: For the steppings affected, see the *Summary Table of Changes*.

A79. VMLAUNCH/VMRESUME may not fail when VMCS is programmed to cause VM exit to return to a different mode

Problem: VMLAUNCH/VMRESUME instructions may not fail if the value of the “host address-space size” VM exit control differs from the setting of IA32_EFER.LMA.

Implication: Programming the VMCS to allow the monitor to be in different modes prior to VMLAUNCH/VMRESUME and after VM exit may result in undefined behavior.

Workaround: Software should ensure that “host address-space size” VM exit control has the same value as IA32_EFER.LMA at the time of VMLAUNCH/VMRESUME.

Status: For the steppings affected, see the *Summary Table of Changes*.

Specification Changes

There are no new Specification Changes for this revision.

The Specification Changes listed in this section apply to the following documents:

1. *Dual-Core Intel® Xeon® Processor 7000 Sequence Datasheet*, Revision 1.0 (Document Number 309626)

All Specification Changes will be incorporated into a future version of the appropriate Dual-Core Intel Xeon processor 7000 sequence processor documentation.

Specification Clarifications

There are no new Specification Clarifications for this revision.

The Specification Clarifications listed in this section apply to the following documents:

1. *Dual-Core Intel® Xeon® Processor 7000 Sequence Datasheet*, Revision 1.0 (Document Number 309626)

All Specification Clarifications will be incorporated into a future version of the appropriate Dual-Core Intel Xeon processor 7000 sequence documentation.

Documentation Changes

Note: Documentation changes for *IA-32 Intel® Architecture Software Developer's Manual*, Volumes 1, 2A, 2B, and 3 will be posted in the separate document *IA-32 Intel® Architecture Software Developer's Manual Documentation Changes*. Follow the link below to become familiar with this file.

<http://developer.intel.com/design/pentium4/specupdt/252046.htm>

There are no new Documentation Changes for this revision.

The Documentation Changes listed in this section apply to the following documents:

1. *Dual-Core Intel® Xeon® Processor 7000 Sequence Datasheet*, Revision 1.0 (Document Number 309626)

All Documentation Changes will be incorporated into a future version of the appropriate Dual-Core Intel Xeon processor 7000 sequence documentation.